Optimal Power Control for CDMA Systems in the Wideband Limit

Tim Holliday, Andrea Goldsmith, and Peter Glynn
Stanford University

Abstract—**Traditional methods for solving multiuser control problems in CDMA systems do not scale well as the number of users in the system increases. Indeed, the size and complexity of the control problem generally grows exponentially with the number of users. In order to find a tractable solution to these problems we develop a wideband limit approximation for dynamic programming problems in CDMA systems. This approximation greatly simplifies the state space and evolution equations for the users in a CDMA cell. The simplified evolution dynamics can then be used to construct a dynamic program that solves for the optimal wideband power control. This wideband convergence result holds for a number of different linear multiuser receiver structures including the MMSE and matched filter receivers.**

I. Introduction

Adaptive power control for wireless systems is a well studied problem. Recently several authors have published power control formulations for a wide variety of applications including joint source/channel coding, delay constrained traffic, delayed channel estimates, and so forth [1], [2], [4], [5]. Typically these formulations involve the solution of a standard stochastic control (e.g. dynamic programming) problem. In many cases, dynamic programming formulations are particularly well suited to the power control problem, since they offer a general modeling and performance constraint framework. However, nearly all of these formulations require a fundamental simplifying assumption about the nature of the interference in a wireless system. That is, the interference observed by any user in the system does not react to the adaptive control chosen by the user. From a modeling standpoint this means a user could select a high transmission power and not expect the observed interference from all the other users in the system to increase as a result. Unfortunately, it is well known that the assumption of unresponsive interference is inaccurate for code-division (CDMA) systems, which we consider in this extended abstract.

The main motivation behind this simplified interference model is to reduce problem complexity to a tractable level. In general, standard dynamic programming formulations have two different methods available to model multiple interacting controllers (i.e. multiple users in a CDMA cell). The first method requires a joint model of the state space and transition dynamics for every user in the system. This type of formulation can only be used effectively for systems with a small number of users. As the number of users increases the size of the state space grows exponentially and the well known "Curse of Dimensionality" precludes a scalable solution. The second method attempts to solve the problem from the perspective of a single user that operates in a system with other responsive users. This type of solution requires a model for the behavior of the other users in the system as a function of the single user's behavior. Unfortunately this solution is also intractable due to the dimensionality problem discussed above. Furthermore, we also need to know the control policy for the users in the system in order to characterize their responsive behavior. However, we will not know the users' control policies until we solve the control problem.

In order to address the difficulties described above, many authors have contributed to a large body of theory on fluid and Brownian approximations for stochastic control problems. These formulations approximate a dynamic system with a scaled version of the system and then solve an optimal control problem for the scaled version. For example, in the Brownian case the system state is scaled by $\sqrt{n}$ and time is scaled by n, where n is a parameter that tends to infinity. As a result, the scaled system characterizes the behavior of the original system over relatively large time scales. The main advantage of this sort of approximation is that the state space of the scaled system collapses to one of substantially lower dimension than the original system. Thus an approximate solution to a previously intractable problem can be found. See [10] for a recently published application of a Brownian approximation in the wireless context.

However, these types of approximations face a unique difficulty when applied to wireless systems. Due to the time scaling, much of the short time scale transitions of the wireless channel are lost. For example, the long time scale behavior of the channels for a collection of 40km/h and 3km/h mobile users might appear to be the same. Hence the optimal control for the scaled version of the system will not account for this difference in user velocity. While this point will often be irrelevant for throughput constrained system analysis (e.g. this scaling works quite well for capacity problems), it is a critical issue in the analysis of systems with delay sensitive traffic [4]. Indeed, Brownian approximations for systems with delay constrained traffic are only effective when the delay tolerance of the traffic is quite high [9]. Our goal is to find a limit regime that achieves a reduction in complexity similar to the approximations discussed above without scaling time and losing channel resolution.

One option that has not been explored previously is the second solution method described above, where the system is modeled from the perspective of a single user. In this model we have another limiting regime to consider that is particular well suited to CDMA systems, namely the wideband limit. In this regime, the original system is approximated by a scaled version, where the scaling is with respect to frequency (i.e. the number of chips in the spreading sequence) and the number of users. The main advantage of this limit approximation over the Brownian model is that we do not lose any time resolution of the wireless channel between each user and the base station, and can therefore evaluate system performance for delay constrained traffic.

Several authors have approximated the signal to interference ratio (SIR) for a variety of receiver structures [12] in the wideband limit. In this paper, rather than approximate the SIR in a single timeslot, our goal is to find a wideband approximation for the state evolution equations of the users in a CDMA cell. As in [12], we will show that the appropriate state description for a CDMA cell is the empirical distribution for all of the users in the cell. We then show that the evolution equation for the cell empirical distribution converges to a deterministic non-linear mapping in the wideband limit. We then use this approximation for the evolution dynamics to construct and solve an optimal control problem for the wideband system. As noted in [12], the wideband approximation can be applied to a variety of receiver structures. Along with our wideband approximation, we will discuss the different optimal control policies that arise from different receiver structures, performance metrics, and performance constraints.

The rest of this paper is organized as follows. In the next section we present the system model and the components required to construct the dynamic programming problem. In the third section we develop the wideband limit approximation and show how it allows us to greatly reduce the complexity of our dynamic program. We conclude in Section IV.

II. System Model and Problem Formulation

In this section we develop a model for a wireless user and a CDMA cell. Our goal is to construct a state space model and evolution equations for each user and then use these components to construct a model for an entire cell. We then use this cell model to formulate an dynamic programming problem. Here we will consider a fairly simple model for the mobile devices and their wireless channels. Though far more general link adaptation strategies and models are possible (e.g. see [4], [5]), our goal is to demonstrate the extension of previous research to account for responsive interference rather than a discussion of complicated system models.

A. The Data Traffic and Buffer Model

Consider a group of K wireless users transmitting data to a base station. In each timeslot every user may generate a packet with probability λ. When packet j is generated at time t_j it announces a deadline d_j such that if it is not transmitted by time $e_j = t_j + d_j$ (the extinction time) it is dropped from the data buffer and considered lost. We assume all packets have a common deadline $d_j = d$. Therefore we can assume each user has a finite size data buffer capable of storing d time slots worth of data. Let $y_i(t)$ be the state of the ith user's data buffer at time t.

B. The Wireless Channel Model

We assume the variations in the wireless channel gain between each user and the base station can be described by a finite-state discrete-time Markov chain (DTMC). (See [3], [7] for example channel models.) Further assume each user's channel follows an independent copy of the same DTMC. We will denote the state of the ith user's channel gain at time t as $z_i(t)$.

C. The State Space Model

We define the state of user i at time t as $x_i(t) = (y_i(t), z_i(t)) \in \mathcal{X}$, where the set $\mathcal{X}$ consists of all combinations of buffer states and channel states for a single user. Using the state of each user we can define a state for the entire system as the empirical distribution $\Pi_K(t)$ over all K users in the cell at time t:

$$\Pi_K(j,t) = \frac{1}{K}\sum_{k=1}^{K} I_{[x_k(t)=j]}, \tag{1}$$

where $I_{[x_k(t)=j]} = 1$ if mobile k is in state j at time t and zero otherwise (i.e. $I_{[A]}$ is the standard indicator random variable). We should point out that for finite K the empirical distribution $\Pi_K(t)$ is a random vector.

D. Power Control

In each time slot every mobile selects a power level $a_i(t) \in A$, where A is a finite set of transmission powers. We assume each mobile chooses a transmission power according to a power control policy $g\left(x_i(t), \Pi_K(t)\right)$, which maps the state of the ith user and the empirical distribution $\Pi_K(t)$ into a probability distribution across A. This determines the probability that the ith user takes an action $a_i(t) \in A$ at time t (i.e. we are allowing for randomized controls). Finally, we define a cost vector $r(a)$, which determines the cost of choosing action a. In this paper our goal is to minimize average transmission power, hence we set $r(a) = a$.

E. Receiver Structure and SIR

As in [12], we assume the ith user is assigned a N length spreading sequence S_i, which is constructed as follows:

$$S_i = \frac{1}{\sqrt{N}}\{(V_{i1}, \ldots, V_{iN}\}, \quad i = 1, \ldots, K, \tag{2}$$

where the elements V_{ij} are independently and identically distributed with mean zero, variance one, and $E[V_{ij}^4] < \infty$.

The structure of the receiver at the base station is a linear demodulator. Therefore the received SIR for user i at time t with K users in the system is

$$SIR_K(i,t) = \frac{(c_i^T S_i)^2 a_i(t) z_i(t)}{(c_i^T c_i)\sigma^2 + \sum_{j=1 j\neq i}^{K} (c_i^T S_i)^2 a_j(t) z_j(t)}, \tag{3}$$

where σ^2 is the background noise power, $a_i(t)z_i(t)$ is the received power from user i, and c_i is the demodulator for user i. The traditional CDMA approach simply chooses $c_i = S_i$ (i.e. a matched filter receiver). The optimal MMSE receiver chooses c_i to maximize the received signal-to-interference ratio $SIR_K(i,t)$, see [8] for details.

The above linear demodulator provides a feature that is critical to our analysis. Namely, the wideband limiting SIR for the ith user can be computed directly [12] for many different receiver structures; provided we know the power control policy $g\left(x_i(t), \Pi_K(t)\right)$, the state of the ith user $x_i(t)$, and the empirical distribution of all users $\Pi_K(t)$.

Let the jth component of the vector $h(g,t)$ be the received power from user j at time t given power control g. Then if the receiver is a matched filter, the SIR for the ith user as $K, N \to \infty$ with $\frac{K}{N} \to \alpha$ is

$$SIR_\infty(i,t) = \frac{H_i(g,t)}{\sigma^2 + \alpha\pi(t)h(g,t)}. \tag{4}$$

Likewise, for the MMSE receiver we have

$$SIR_\infty(i,t) = \tag{5}$$

$$\frac{h_i(g,t)h_j(g,t)}{\sigma^2 + \alpha\sum_{j=1,j\neq i}^{K} \pi_j(t)\frac{h_i(g,t)h_j(g,t)}{h_i(g,t)+h_j(g,t)SIR_\infty(i,t)}}. \tag{6}$$

Notice that the SIR calculation for the matched filter receiver is a straightforward computation while the MMSE receiver structure requires us to solve a fixed-point equation for the SIR.

As a technical aside, for these limits to exist we must require that the limiting empirical distribution, $\pi(t)$, exists at time t. That is,

$$\pi(t) = \lim_{K,N\to\infty} \Pi_K(t) \tag{7}$$

must be a deterministic vector. We will discuss this particular issue in Section III.

F. Evolution Equations

Now that we have defined our state space model, power control policy, and receiver structure, we can set up a system of equations describing the evolution of a CDMA cell over time. We assume the probability of mobile i successfully transmitting a packet of data at time t depends only on $SIR_K(i,t)$. Recall that we also assumed the transitions of the ith user's wireless channel are independent of all other users' states and power choices. Therefore the transition probabilities for the ith user are only a function of $x_i(t)$, the power control $g\left(x_i(t), \Pi_K(t)\right)$, and the empirical distribution $\Pi_K(t)$. Hence we define a transition matrix $P\left(g, \Pi_K(t)\right)$ whose jkth element is

$$P\left(g, \Pi_K\right)_{jk} = P\left(x_i(t+1) = k | x_i(t) = j, g, \Pi_K(t)\right) \tag{8}$$

$$= f_k\left(x_i(t), g, \Pi_K(t)\right), \tag{9}$$

where f is a continuous function that maps $x_i(t)$, g, and $\Pi_K(t)$ to a probability distribution on the states in X.

As we mentioned above, the empirical distribution $\Pi_K(t)$ is a random vector, and therefore the transition matrix $P(g, \Pi_K(t))$ is also random. As one might guess, the number of different values $\Pi_K(t)$ can take will grow very quickly as the number of users in the system increases. Furthermore, the probability distribution for $\Pi_K(t)$ can be particularly difficult to compute since it involves the summation of K random variables that are neither independent nor identically distributed.

G. *Dynamic Programming Formulation*

The term "dynamic program" refers to a wide variety of formulations and optimization techniques used to solve optimal control problems. However, one common feature among all of these techniques is the notion of a value function – a quantity that determines the cost or reward associated with a particular control policy (e.g. capacity, throughput, or power consumption). The goal of a dynamic program is to optimize the value function subject to constraints on system behavior (e.g. average power consumption or probability of deadline expiration). In this paper we use an infinite horizon value function, which assigns equal weight to the cost associated with the selected control in each time slot. Hence if $R_i(g,t)$ is the random process describing the cost associated with the control for user i at time t, we write the value of the power control policy g as

$$V_K(g,i) = \lim_{T\to\infty} \frac{1}{T+1} \sum_{t=0}^{T} E_{x,\Pi}[R_i(g,t)], \tag{10}$$

where the expectation is taken with respect to the state $x_i(t)$ and empirical distribution $\Pi_K(t)$.

In order to solve this dynamic program for the problem formulation discussed in the previous section we must find the power control g that minimizes our value function (10). At this point we can see how the curse of dimensionality makes a general solution to this problem intractable. In order to optimize (10) we must be able to evaluate the value function for any choice of control. This requires us to find a stationary distribution of $R_i(g,t)$, which requires knowledge of the stationary distribution of $\Pi_K(t)$. However, as we mentioned above, the number of possible states for $\Pi_K(t)$ can be extremely large. Indeed it is quite easy to create trivial examples with only three or four users such that computing the stationary distribution of $\Pi_K(t)$ is computationally infeasible. In the next section we will show that these complexity issues are largely eliminated when we consider the appropriate wideband limit of $\Pi_K(t)$.

III. A Wideband Approximation

Our goal in this section is to simplify the evolution equations and value function developed above. In particular, we seek an analytical regime that collapses the random transition matrices (8) to deterministic quantities. The simplified transition dynamics will allow us to substantially simplify our state space description of a CDMA cell and optimize our proposed value function. As mentioned above, the appropriate regime is the limit as $N, K \to \infty$ and $\frac{N}{K} \to \alpha$. In order to simplify and optimize our value function (10), we will use the propositions in the following section. (Due to space constraints we will not provide proofs for these propositions, the reader is directed to [6] for more details.)

A. *Transition Matrices and the Empirical Distribution*

We assume that the K mobiles' states are initialized at time $t = 0$ independently according to a probability distribution $\Pi_K(0)$. As in [12], we further assume the following limit exists

$$\lim_{K,N\to\infty} \Pi_K(0) = \pi(0). \tag{11}$$

We assume this condition for the sake of simplicity, it can be relaxed somewhat at the cost of more complicated proof techniques. Then at time $t = 0$, our transition matrix for a single user converges to a deterministic quantity,

$$\lim_{K,N\to\infty} P(g, \Pi_K(0)) = P(g, \pi(0)). \tag{12}$$

The reasoning behind this statement is straightforward. From [12] we know the SIR for each user converges as $K, N \to \infty$ provided the empirical distribution $\pi(0)$ exists. Since our transition matrices for each user are continuous functions of the user's SIR and control policy, the matrices also converge. In order to see how this affects our transition dynamics for time $t > 0$ we consider the following,

Proposition 1: As $K, N \to \infty$ the empirical distribution $\Pi_K(1)$ at time $t = 1$ converges to the deterministic non-linear matrix function,

$$\pi(1) = \pi(0)P(g, \pi(0)), \tag{13}$$

and by induction this equation holds for general time t such that,

$$\lim_{K\to\infty} \Pi_K(t+1) = \pi(t+1) = \pi(t)P(g, \pi(t)). \tag{14}$$

Proof: See [6]

Proposition 2: The mapping $\pi = \pi P(g, \pi)$ has a unique fixed point and there is a unique solution to the equation $\pi(\infty) = \pi(\infty)P(g, \pi(\infty))$. Therefore there is a unique limit $\lim_{t\to\infty} \pi(t) = \pi(\infty)$.

Proof: See [6].

B. *Value Function Simplification*

Now that we have the existence of a unique limit for the empirical distribution of users in a CDMA cell, we can greatly simplify the value function (10).

Proposition 4: The value function proposed in equation (10) simplifies to the vector product $\pi(\infty, g)r(g)^T$, where $\pi(\infty, g)$ is the limit of the empirical distribution as a function of the control policy g.

Proof: Define a mapping $T^n(g,\pi)$ such that $T(g,\pi) = \pi P(g,\pi)$, $T^2(g,\pi) = (\pi P(g,\pi))P(g, \pi P(g,\pi))$, etc., which maps a starting vector $\pi(0)$ n steps into the future. Then our value function as $K, N \to \infty$ is

$$\begin{aligned}
V_\infty(g) &= \lim_{T\to\infty} \lim_{K\to\infty, N\to\infty} \left[\frac{1}{T+1} \sum_{n=0}^{T} E_{X,\Pi}[R(g)] \right] && (15)\\
&= \lim_{T\to\infty} \left[\frac{1}{T+1} \sum_{n=0}^{T} T^n\left(g, X(0)\right) r(g)^T \right] && (16)\\
&= \lim_{T\to\infty} \left[\frac{1}{T+1} \sum_{n=0}^{T} T^n\left(g, X(0)\right) \right] r(g)^T && (17)\\
&= \pi(\infty, g) r(g)^T, && (18)
\end{aligned}$$

where the last equality comes from the existence of a unique fixed point of the mapping T.

C. Optimization Formulation

Now that we have simplified our value function and dynamic equations we can formulate an optimization problem that minimizes our value function subject to constraints on system behavior. We will characterize our performance constraints by limiting the frequency of visits to "undesirable" states (e.g. where a data packet is lost due to deadline expiration). We will represent this constraint (or possibly a set of constraints) through a matrix A and vector b such that our performance constraints are met if $\pi(\infty, g)A \leq b$. With the performance constraint and the development in the previous sections we can use the following non-linear program to solve for the power control policy that minimizes our value function (15):

$$\min_g \pi(\infty, g) r(g)^T \tag{19}$$

subject to:

$$\begin{aligned} \pi(\infty, g) &= \pi(\infty, g) P\left(\pi(\infty, g)\right) \\ \pi(\infty, g)\mathbf{1}^T &= 1 \\ \pi(\infty, g)A &\leq b \\ \pi(\infty, g) &\geq 0 \end{aligned}$$

Intuitively, we are minimizing the value function subject to the constraint that the steady-state empirical distribution $\pi(\infty, g)$ satisfies the fixed point equation (20). As well, $\pi(\infty, g)$ must be a probability distribution and satisfy our performance constraints.

The above non-linear program is very similar in appearance to the linear program that would be derived if the users in the system did not respond to each others power choice. In the non-responsive interference case, the transition matrix $P(\pi(\infty, g))$ in the first constraint would simply be a deterministic matrix P, rather than a function of the empirical distribution and power control. Hence, we have reduced the size of our multiuser problem to that of a single user problem, albeit with non-linear dynamics.

D. Technical Details

In the previous sections we did not discuss a few key technical questions. First, can we always find a solution to our proposed non-linear program? Second, what are the drawbacks to applying the wideband approximation to a finite bandwidth system?

Unfortunately, it is possible to show that our proposed non-linear program is not convex. Clearly this is a significant drawback to our formulation since we are not guaranteed to find the optimal power control policy when we solve the optimization problem (19). However, one can show that it is possible to get arbitrarily close to the optimal solution by solving a sequence of linear programs. A detailed discussion of this proof is reserved for [6].

The answer to the second question requires two parts, addressing both the theoretical and practical aspects of our approximation. The critical theoretical question is whether or not our exchange of limits in (15) was valid. In our limiting value function $V_\infty(g)$ we first took the limit as $K \to \infty$ and then computed the value function using the dynamics for this limiting system. From a theoretical standpoint it would be reassuring if we could show that the value functions for the finite K systems, $V_K(g)$, converge to $V_\infty(g)$. Though we cannot compute the finite K value functions directly due to the complexity issues mentioned above, we can indeed show that $V_K(g) \to V_\infty(g)$. Once again, a detailed proof is reserved for [6]. The practical issue is the similarity in performance of a finite user system to the theoretical optimum of the wideband system. This concern will be addressed at the conference by comparing the theoretical wideband results to finite bandwidth simulations.

E. Numerical Example

Due to space constraints we cannot present a detailed numerical analysis section in this extended abstract. In [6] and at the conference we will included a detailed example comparing the optimal control for matched filter and MMSE receivers in the wideband limit. In particular, we will examine the minimum average power consumption with constraints on traffic delay versus system load (i.e. the load parameter α). As well, we will consider a throughput maximization problem with constraints on delay and power consumption. Finally, we will present simulation results that evaluate the performance of the wideband control policies in systems with finite bandwidth and finite numbers of users. The results will show that the wideband approximation works well for the bandwidth specifications of typical wideband systems.

IV. Conclusions

We have presented a wideband limit approximation for dynamic programming problems in CDMA systems. In traditional dynamic programming formulations the size and complexity of the problem grows exponentially with the number of users. However, in the wideband limit we showed that the optimal control problem reduces to the size of a single-user control problem with non-linear dynamics. This state-space collapse is somewhat analogous to the complexity reduction that can be achieved by standard fluid or Brownian approximations. However, our formulation provides a key benefit over these standard techniques in that we do not sacrifice any time scale resolution of the wireless channel. Hence the wideband limit provides an appropriate regime for evaluating the performance of CDMA systems with delay sensitive traffic.

References

[1] N. Bambos, S. Kandukuri, *Power Controlled Multiple Access in Wireless Communications Networks*, IEEE INFOCOM 2000.

[2] S. Gitzenis, N. Bambos, Power-Controlled Data Prefetching/Caching in Wireless Packet Networks, IEEE INFOCOM 2002.

[3] M. Gunmundson, Correlation Model for Shadow Fading in Mobile Radio Systems, Electronics Letters, Vol. 27, November 1991.

[4] T. Holliday, A. Goldsmith, and P. Glynn, Optimal Power Control and Source-Channel Coding for Delay Constrained Traffic over Wireless Channels, IEEE ICC 2002.

[5] T. Holliday, A. Goldsmith, and P. Glynn, Wireless Link Adaptation Policies: QoS for Deadline Constrained Traffic with Imperfect Channel Estimates, IEEE ICC 2002

[6] T. Holliday, A. Goldsmith, and P. Glynn, A Wideband Limit Approximation for Optimal Power Control in CDMA Systems, Journal version in preparation.

[7] Y.Y. Kim, S. Li, *Modeling Fast Fading Channel Dynamics for Packet Data Performance Analysis*, Proceedings IEEE INFOCOM '98.

[8] U. Madhow and M. Honig, MMSE Interference Suppression for Direct-Sequence Spread-Spectrum CDMA, IEEE Transactions on Communication, vol. 42, pp.3178-3188, Dec. 1994.

[9] M. Rubino, M. Harrison, Mananging Workflow in a Flexible Server System with Delay Constraints, Preprint, October 2002.

[10] A.L. Stolyar, MaxWeight scheduling in a generalized switch: State space collapse and equivalent workload minimization under complete resource pooling, Preprint, September 2001.

[11] G. Stuber, *Principles of Mobile Communication*, Kluwer Academic Publishers, 1999.

[12] D. Tse and S. Hanly, Linear Multiuser Receivers: Effective Interference, Effective Bandwidth and User Capacity, IEEE Transactions on Information Theory, v.45, No. 2, Mar. 1999, pp. 641-657.